

TITLE OF THE INVENTION

ECHO CANCELLER EMPLOYING DUAL-H ARCHITECTURE
HAVING SPLIT ADAPTIVE GAIN SETTINGS

CROSS-REFERENCE TO RELATED APPLICATIONS

The following applications, filed on even date, herewith, are incorporated by reference: USSN _____. (Attorney Docket No. 11724US01), "Echo Canceller Employing Dual-H Architecture Having Improved Coefficient Transfer"; USSN _____, (Attorney Docket No. 11998US01), "Echo Canceller Employing Dual-H Architecture Having Improved Double-Talk Detection"; USSN _____, (Attorney Docket No. 1999US01), "Echo Canceller Employing Dual-H Architecture Having Improved Non-Linear Echo Path Detection"; USSN _____, (Attorney Docket No. 12000US01), "Echo Canceller Employing Dual-H Architecture Having Variable Adaptive Gain Settings"; USSN _____. (Attorney Docket No. 12001US01), "Echo Canceller Employing Dual-H Architecture Having Improved Non-Linear Processor."

STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH
OR DEVELOPMENT

Not Applicable

BACKGROUND OF THE INVENTION

Long distance telephone facilities usually comprise four-wire transmission circuits between switching offices in different local exchange areas, and two-wire circuits within each area connecting individual subscribers with the switching office. A call between subscribers in different exchange areas is carried over a two-wire circuit in each of the areas and a four-wire circuit between the areas, with conversion of speech energy between the two and four-wire circuits being effected by hybrid circuits. Ideally, the hybrid circuit input ports perfectly match the impedances of the two and four-wire circuits, and its balanced network impedance perfectly matches the impedance of the two-wire circuit. In this manner, the signals transmitted from one exchange area to the other will not be reflected or returned to the one area as echo. Unfortunately, due to impedance differences which inherently exist between different two and four-wire circuits, and because impedances must be matched at each frequency in the voice band, it is virtually impossible for a given hybrid circuit to perfectly match the impedances of any particular two and four-wire transmission circuit. Echo is, therefore, characteristically part of a long distance telephone system.

Although undesirable, echo is tolerable in a telephone system so long as the time delay in the echo path is relatively short, for example, shorter than about 40 milliseconds. However, longer echo delays can be distracting or utterly confusing to a far end speaker, and to reduce the same to a tolerable level an echo canceller may be used toward each end of the path to cancel echo which otherwise would return to

the far end speaker. As is known, echo cancellers monitor the signals on the receive channel of a four-wire circuit and generate estimates of the actual echoes expected to return over the transmit channel. The echo estimates are then applied to a subtractor circuit in the transmit channel to remove or at least reduce the actual echo.

In simplest form, generation of an echo estimate comprises obtaining individual samples of the signal on the receive channel, convolving the samples with the impulse response of the system and then subtracting, at the appropriate time, the resulting products or echo estimates from the actual echo on the transmit channel. In actual practice generation of an echo estimate is not nearly so straightforward.

Transmission circuits, except those which are purely resistive, exhibit an impulse response has amplitude and phase dispersive characteristics that are frequency dependent, since phase shift and amplitude attenuation vary with frequency. To this end, a suitable known technique for generating an echo estimate contemplates manipulating representations of a plurality of samples of signals which cause the echo and samples of impulse responses of the system through a convolution process to obtain an echo estimate which reasonably represents the actual echo expected on the echo path. One such system is illustrated in FIG. 1.

In the system illustrated in FIG. 1, a far end signal x from a remote telephone system is received locally at line 10. As a result of the previously noted imperfections in the local system, a portion of the signal x is echoed back to the remote site at line 15 along with the signal v from the local telephone system. The

echo response is illustrated here as a signal s corresponding to the following equation:

where h is the impulse response of the echo characteristics. As such, the signal sent from the near end to the far end, absent echo cancellation, is the signal y , which is the sum of the telephone signal v and the echo signal s . This signal is illustrated as y at line 15 of FIG. 1.

To reduce and/or eliminate the echo signal component s from the signal y , the system of FIG. 1 uses an echo canceller having an impulse response filter \bar{h} that is the estimate of the impulse echo response h . As such, a further signal \bar{s} representing an estimate of echo signal s is generated by the echo canceller in accordance with the following equation:

$$\bar{s} = \bar{h} * x$$

The echo canceller subtracts the echo estimate signal \bar{s} from the signal y to generate a signal e at line 20 that is returned to the far end telephone system. The signal e thus corresponds to the following equation:

$$e = s + v - \bar{s} \approx v$$

As such, the signal returned to the far end station is dominated by the signal v of the near end telephone system. As the echo impulse response \bar{h} more closely

correlates to the actual echo response h , then \bar{s} more closely approximates s and thus the magnitude of the echo signal component s on the signal e is more substantially reduced.

The echo impulse response model \bar{h} may be replaced by an adaptive digital filter having an impulse response \hat{h} . Generally, the tap coefficients for such an adaptive response filter are found using a technique known as Normalized Least Mean Squares adaptation.

Although such an adaptive echo canceller architecture provides the echo canceller with the ability to readily adapt to changes in the echo path response h , it is highly susceptible to generating sub-optimal echo cancellation responses in the presence of "double talk" (a condition that occurs when both the speaker at the far end and the speaker at the near end are speaking concurrently as determined from the viewpoint of the echo canceller).

To reduce this sensitivity to double-talk conditions, it has been suggested to use both a non-adaptive response and an adaptive response filter in a single echo canceller. One such echo canceller is described in USPN 3,787,645, issued to Ochiai et al on January 22, 1974. Such an echo canceller is now commonly referred to as a dual-H echo canceller.

Although the dual-H echo canceller architecture of the '645 patent provides substantial improvements over the use of a single filter response architecture, the

'645 patent is deficient in many respects and lacks certain teachings for optimizing the use of such a dual-H architecture in a practical echo canceller system. For example, the present inventors have recognize that the adaptation gain used to adapt the tap coefficients of the adaptive filter may need to be altered based on certain detected conditions. These conditions include conditions such as double-talk, non-linear echo response paths, high background noise conditions, etc.. The present inventors have recognized the problems associated with the foregoing dual-H architecture and have provided solutions to such conditions.

BRIEF SUMMARY OF THE INVENTION

An echo canceller circuit is set forth. The echo canceller circuit includes a digital filter having adaptive tap coefficients to simulate an echo response occurring during a call. The adaptive tap coefficients are updated during the call using a Mean Squares or at least squares process. A tap energy detector is also employed. The tap energy detector identifies and divides groups of taps having high energy from groups of taps having low energy. The high energy tap groups are generally smaller in number than the low energy tap groups. The high energy tap groups are adapted separately from the low energy tap groups using the Mean Squares process. Still further, the high energy tap groups may be adapted using an adaptive gain constant a while the low energy tap groups are adapted using an adaptive gain constant a' , wherein $a > a'$.

BRIEF DESCRIPTION OF THE SEVERAL
VIEWS OF THE DRAWINGS

Figure 1 is a block diagram of a conventional canceller.

Figure 2 is a schematic block diagram of an echo canceller that operates in accordance with one embodiment of the present invention.

Figure 3 is a flow chart illustrating one manner of carrying out coefficient transfers wherein the transfer conditions may be used to implement double-talk detection in accordance with one embodiment of the present invention.

Figure 4 is a flow chart illustrating a further manner of carrying out coefficient wherein the transfer conditions may be used to implement the double-talk detection in accordance with one embodiment of the present invention.

Figure 5 illustrates an exemplary solution surface for the adaptive filter whereby the desired result is achieved at the solution matching the echo response of the channel.

Figure 6 illustrates one manner of checking for various echo canceller conditions and responding to these conditions using a change in the adaptive gain setting of the adaptive filter of the echo canceller.

Figure 7 illustrates a typical linear echo path response for the \hat{h} filter and one manner of identifying and separating one or more groups of high energy taps from one or more groups of low energy taps for separate adaptation.

Figure 8 illustrates one manner of implementing an echo canceller system employing the present invention.

DETAILED DESCRIPTION OF THE INVENTION

Figure 2 illustrates one embodiment of a dual-h echo canceller suitable for use in implementing the present invention. As illustrated, the echo canceller, shown generally at 25, includes both a non-adaptive filter \bar{h} and an adaptive filter \hat{h} to model the echo response h . Each of the filters \bar{h} and \hat{h} are preferably implemented as digital filters, such as finite impulse response (FIR) filters comprising a plurality of taps each having a corresponding tap coefficient. This concept may be extended to IIR filters as well. If FIR filters are used, the duration of each of the FIR filters should be sufficient to cover the duration of the echo response of the channel in which the echo canceller 25 is disposed.

The output of the non-adaptive filter \bar{h} is available at the line 30 while the output of the adaptive filter \hat{h} is available at line 35. Each of the signals at lines 30 and 35 are subtracted from the signal-plus-echo signal of line 40 to generate echo compensated signals at lines 50 and 55, respectively. A switch 45, preferably a software switch, may be used to selectively provide either the output signal at the line 50 or the output signal at line 55 to the echo canceller output at line 60. The switch 45 may be used to provide the echo compensation based on the \hat{h} filter during convergence and then be switched to provide the echo compensation based on the \bar{h} filter after convergence. Further, the switch 45 is directed to provide the echo compensation based on the \bar{h} filter in response to the detection of a

double-talk condition.

A transfer controller 65 is used to transfer the tap coefficients of filter \hat{h} to replace the tap coefficients of filter \bar{h} . As illustrated, the transfer controller 65 is connected to receive a number of system input signals. Of particular import with respect to the present invention, the transfer controller 65 receives the signal-plus-echo response y and each of the echo canceller signals \bar{e} and \hat{e} at lines 50 and 55, respectively. The transfer controller 65 is preferably implemented in the software of one or more digital signal processors used to implement the echo canceller 25.

As noted above, the art is substantially deficient of teachings with respect to the manner in which and conditions under which a transfer of tap coefficients from \hat{h} to \bar{h} is to occur. The present inventors have implemented a new process and, as such, a new echo canceller in which tap coefficient transfers are only made by the transfer controller 65 when selected criterion are met. The resulting echo canceller 25 has substantial advantages with respect to reduced double-talk sensitivity and increased double-talk detection capability. Further, it ensures a monotonic improvement in the estimates \bar{h} .

The foregoing system uses a parameter known as echo-return-loss-enhancement (ERLE) to measure and keep track of system performance. Two ERLE parameter values are used in the determination as to whether the transfer controller 65 transfers the tap coefficients from \hat{h} to \bar{h} . The first parameter, \bar{E} , is defined in

the following manner:

$$\bar{E} = \frac{y}{\bar{e}}$$

Similarly, the parameter \hat{E} is defined as follows:

$$\hat{E} = \frac{y}{\hat{e}}$$

Each of the values \hat{E} and \bar{E} may also be averaged over a predetermined number of samples to arrive at averaged \hat{E} and \bar{E} values used in the system for the transfer determinations.

Figure 3 illustrates one manner of implementing the echo canceller 25 using the parameters \hat{E} and \bar{E} to control tap coefficients transfers between filter \hat{h} to \bar{h} . As illustrated, the echo canceller 25 provides a default \bar{h} set of coefficients at step 80 during the initial portions of the call. After the tap coefficients values for \bar{h} have been set, a measure of \bar{E} is made at step 85 to measure the performance of the tap coefficient values of filter \bar{h} .

After the initialization sequence of steps 80 and 85, or concurrent therewith, the echo canceller 25 begins and continues to adapt the coefficients of \hat{h} to more adequately match the echo response h of the overall system. As noted in Figure 3, this operation occurs at step 90. Preferably, the adaptation is made using a Normalized Least Mean Squares method, although other adaptive methods may also

be used (e.g., LMS and RLS).

After a period of time has elapsed, preferably, a predetermined minimum period of time, the echo canceller 25 makes a measure of \hat{E} at step 95. Preferably, this measurement is an averaged measurement. At step 100, the echo canceller 25 compares the value of \hat{E} with the value of \bar{E} . If the value of \hat{E} is greater than the value of \bar{E} , the tap coefficients of filter \hat{h} are transferred to replace the tap coefficients of filter \bar{h} at step 105. If this criterion is not met, however, the echo canceller 25 will continue to adapt the coefficients of the adaptive filter \hat{h} at step 90, periodically measure the value of \hat{E} at step 95, and make the comparison of step 100 until the condition is met.

Although not illustrated, other transfer conditions may be imposed in addition to the foregoing. For example, the echo canceller may impose a requirement that a far end signal exist before a transfer may occur. Additionally, transfers may be inhibited during a double-talk condition. Further conditions may also be imposed based on system requirements.

If the echo canceller 25 finds that \hat{E} is greater than \bar{E} , the above-noted transfer takes place. Additionally, the echo canceller 25 stores the value of \hat{E} as a value E_{\max} . This operation is depicted at step 110 of the Figure 3. The value of E_{\max} is thus the maximum value of ERLE that occurs over the duration of the call and at which a transfer has taken place. This further value is used thereafter, in

addition to the \hat{E} and \bar{E} comparison, to control whether the tap coefficients of \hat{h} are transferred by the transfer controller 65 to replace the tap coefficients of \bar{h} . This further process is illustrated that steps 115, 120, and 125 of Figure 3. In each instance, the tap coefficient transfer only occurs when both of the following two conditions are met: 1) \hat{E} is greater than the current \bar{E} , and 2) \hat{E} is greater than any previous value of \bar{E} used during the course of the call. (\hat{E} is greater than E_{\max}). Each time that both criteria are met, the transfer controller 65 of echo canceller 25 executes the tap coefficient transfer and replaces the previous E_{\max} value with the current \hat{E} value for future comparison.

Requiring that \hat{E} be greater than any \bar{E} value used over the course of the call before the coefficient transfer takes place has two beneficial and desirable effects. First, each transfer is likely to replace the prior tap coefficients of filter \bar{h} with a better estimate of the echo path response. Second, this transfer requirement increases the double-talk protection of the echo canceller system. Although it is possible to have positive ERLE \hat{E} during double-talk, the probability that \hat{E} is greater than E_{\max} during double-talk decreases as the value of E_{\max} increases. Thus an undesirable coefficient transfer during double-talk becomes increasingly unlikely as the value of E_{\max} increases throughout the duration of the call.

The echo canceller 25 may impose both an upper boundary and a lower boundary on the value of E_{\max} . For example, E_{\max} may have a lower bounded

value of 6 dB and an upper bounded value of 24 dB. The purpose of the lower bound is to prevent normal transfers during double-talk conditions. It has been shown in simulations using speech inputs that during double-talk, a value of greater than 6 dB ERLE was a very low probability event, thus making it an appropriate value for the initial value of E_{\max} . The upper bound on E_{\max} is used to prevent a spuriously high measurement from setting E_{\max} to a value at which further transfers become impossible.

The value of E_{\max} should be set to, for example, the lower bound value at the beginning of each call. Failure to do so will prevent tap coefficient transfers on a new call until the echo cancellation response of the echo canceller 25 on the new call surpasses the quality of the response existing at the end of the prior call. However, this criterion may never be met during the subsequent call. Thereby causing the echo canceller 25 to operate using sub-optimal tap coefficients values. Resetting the E_{\max} value to a lower value increases the likelihood that a tap coefficient transfer will take place and, thereby, assists in ensuring that the \bar{h} filter uses tap coefficients for echo cancellation that more closely correspond to the echo path response of the new call.

One manner of implementing the E_{\max} value change is illustrated in the echo canceller operations flow-chart of Figure 4. When all transfer conditions are met except \hat{E} greater than E_{\max} , and this condition persists for a predetermined

duration of time, the echo canceller 25 will reset the E_{\max} value to, for example, the lower bound value. In the exemplary operations shown in Figure 4, the echo canceller 25 determines whether \hat{E} is greater than the lower bound of E_{\max} at step 140 and less than the current value of E_{\max} at step 145. If both of these conditions remain true for a predetermined period of time as determined at step 150, and all other transfer criteria have been met, the echo canceller 25 resets the E_{\max} value to a lower value, for example, the lower bound of the E_{\max} value, at step 155. This lowering of the E_{\max} value increases the likelihood of a subsequent tap coefficient transfer.

Choosing values for the lower and upper bound of E_{\max} other than 6 dB and 24 dB, respectively, is also possible in the present system. Choosing a lower bound of E_{\max} smaller than 6 dB provides for a relatively prompt tap coefficient transfer after a reset operation or a new call, but sacrifices some double-talk protection. A value greater than 6 dB, however, inhibits tap coefficient transfer for a longer period of time, but increases the double-talk immunity of the echo canceller. Similarly, varying the value of the predetermined wait time T before which E_{\max} is reset may also be used to adjust echo canceller performance. A shorter predetermined wait time T produces faster reconvergence transfers, but may sacrifice some double-talk immunity. The opposite is true for larger predetermined wait time values.

A further modification of the foregoing echo canceller system relates to the value stored as E_{\max} at the instant of tap coefficient transfer. Instead of setting E_{\max}

equal to the \hat{E} value at the transfer instant, E_{\max} may be set to a value equal to the value of \hat{E} minus a constant value (e.g., one, three, or 6 dB). At no time, however, should the E_{\max} value be set to a value that is below the lower bound value for E_{\max} . Additionally, a further condition may be imposed in that a new softened E_{\max} is not less than the prior value of E_{\max} . The foregoing "softening" of the E_{\max} value increases the number of transfers that occur and, further, provides more decision-making weight to the condition of \hat{E} being larger than \bar{E} . Further details with respect to the operation of the echo canceller coefficient transfer process are set forth in the co-pending patent application titled "ECHO CANCELLER HAVING THE IMPROVED TAP COEFFICIENT TRANSFER", (Attorney Docket No. ____) filed on even date herewith.

Preferably, the adaptive filter \hat{h} uses a Normalized Least Mean Square (NLMS) adaptation process to update its tap coefficients. In accordance with the process, coefficients are adapted at each time n for each tap $m = 0, 1, \dots, N-1$ in accordance with the following equation:

$$\hat{h}_{n+1}(m) = \hat{h}_n(m) + \frac{a_n}{\sum_{i=0}^{N-1} x_i^2} e_n x_{n-m} \quad \text{for } m=0,1,\dots,N-1$$

where $\hat{h}_n(m)$ is the m^{th} tap of the echo canceller, x_n is the far-end signal at time n , e_n is the adaptation error of time n , and a_n is the adaptation gain at time n .

It is also possible to use the NLMS adaptive process can to adapt coefficients a D.C. tap, h_{dc} . If so desired, the above equation in may be modified to the following:

$$\hat{h}_{n+1}(m) = \hat{h}_n(m) + \frac{a_n}{\sum_{i=0}^{N-1} x_i^2 + x_{dc}^2} e_n x_{n-m} \text{ for } m=0,1,\dots,N-1$$

Additionally, the D.C. tap may be adapted in accordance with the following equation:

$$\hat{h}_{dc,n+1}(m) = \hat{h}_{dc,n}(m) + \frac{a_n}{\sum_{i=0}^{N-1} x_i^2 + x_{dc}^2} e_n x_{dc} \text{ for } m=0,1,\dots,N-1$$

where x_{dc} is a constant.

The foregoing adaptation process will converge in the mean-square sense to the correct solution the echo path response h if $0 < a_n < 2$. Fastest convergence occurs when $a = 1$. However, for $0 < a \leq 1$, the speed of convergence to h is traded-off against steady-state performance.

Figure 5 is provided to conceptualize the effect of the adaptation gain on the filter response. The graph of Figure 5 includes an error performance surface 185 defined to be the mean square error between \hat{h} and h , to be a $N+1$ dimensional bowl. In the present case, $N = 2$. Each (x,y) point in the bowl is connected to a value z which corresponds to the mean-square error for each corresponding (x,y) value \hat{h} (of length N). The bottom of the bowl is the \hat{h} which produces the least mean-square error, i.e. h . The NLMS process iteratively moves the \hat{h} towards h at the bottom of the performance surface as shown by arrow 190. When $a = 1$, \hat{h} moves to the bottom of the bowl most quickly, but once the bottom is reached, the adaptation process continues to bounce \hat{h} around the true h bottom of the bowl, i.e. $E[\hat{h}] = h$ but $\hat{h} \neq h$. If a small a is used, then the steady-state error is smaller (\hat{h} will remain closer to h), but \hat{h} requires a longer time to descend to the bottom of the bowl, as each step is smaller.

In some cases, as the present inventors have recognized, the performance surface will temporarily change. In such situations, it becomes desirable to suppress the \hat{h} from following these changes. This presents a challenge to choose the best a for each scenario.

Figure 6 illustrates operation of the echo canceller 25 in response to various detected scenarios. It will be recognized that the sequence of detecting the various

conditions that is set forth in Figure 6 is merely illustrative and may be significantly varied. Further, it will be recognized that the detection and response to each scenario may be performed concurrently with other echo canceller processes. Still further, it will be recognized that certain detected scenarios and their corresponding responses may be omitted.

In the embodiment of Figure 6, the echo canceller 25 entertains whether or not a double-talk condition exists at step 200. Double talk, as noted above, is defined as the situation when both far-end and near-end talkers speak at the same time during a call. In such a scenario, the adaptive error signal is so severely corrupted by the near-end speaker that it is rendered useless. As such, if a double-talk condition is detected, the echo canceller 25 responds by freezing the adaptation process at step 205, i.e., set $a = 0$, until the double talk ceases.

There are several methods that the echo canceller 25 can use for detecting a double-talk condition. One is to compare the power of the near-end signal to the far-end signal. If the near-end power comes close enough to the far-end power ("close enough" can be determined by the system designer, e.g. within 0 or 6 or 10dB), then double talk can be declared. Another method is to compare the point-by-point magnitudes of the near-end and far-end signals. This search can compare the current $|x|$ with the current $|y|$, the current $|x|$ with the last several $|y|$, the current $|y|$ with the last several $|x|$, etc. In each case, the max $|x|$ and $|y|$ over

the searched regions are compared. If

$$\frac{\max |y|}{\max |x|} > \text{Double Talk Threshold}$$

where $\max |x|$ indicates the maximum $|x|$ over the search region ($|y|$ is similarly defined), then a double-talk condition is declared.

A still further manner of detecting a double-talk condition is set forth in U.S.S.N. ___, titled ___ (Attorney Docket No. ___), the teachings of which are hereby incorporated by reference. As set forth in that patent application, a double-talk condition is declared based on certain monitored filter performance parameters.

It may be possible to further condition the double-talk declaration with other measurements. For example, the current Echo Return Loss (ERL) may be used to set the *Double Talk Threshold* noted above herein. The short-term power of either the far-end, the near-end, or both, may also be monitored to ensure that they are larger than some absolute threshold (e.g. -50dBm or -40dBm). In this manner, a double-talk condition is not needlessly declared when neither end is speaking.

Once a double-talk condition is declared, it may be desirable to maintain the double-talk declaration for a set period time after the double talk condition is met.

Examples might be 32, 64, or 96 msec. After the double-talk condition ceases to exist, the adaptive gain value may be returned to the value that existed prior to the detection of the double-talk condition, or to a predetermined return value.

At step 210, the echo canceller 25 determines whether a high background noise condition is present. A low level of constant background noise can enter from the near-end, for example, if the near-end caller is in an automobile or an airport. Its effects are in some ways similar to that of double-talk, as the near-end double-talk corrupts the adaptive error signal. The difference is that, unlike double talk, near-end background noise is frequently constant, thus setting $\alpha = 0$ until the noise ends is not particularly advantageous. Also background noise is usually of lower power than double-talk. As such, it corrupts the adaptation process but does not render the resulting adapted coefficients unusable.

As illustrated at step 215, it is desirable to choose a gain $0 < \alpha < 1$, i.e. lower the gain from its fastest value of 1 when a high background noise condition is present. While this will slow the adaptation time, the steady state performance increases since the effects of noise-induced perturbations will be reduced. In other words, the tap variance noise is reduced by lowering the adaptation gain α .

Preferably, the background noise is measured as a long-term measurement of the power of when the far-end is silent. As this measurement increases, α decreases. One schedule for setting the adaptive gain α as a function of

background noise level is set forth below.

<u>Background Noise (dBm)</u>	<u>a</u>
> -48	.125
$> -54 \geq -48$.25
$> -60 \geq -54$.5
< -60	1

It will be readily recognized that there are other schedules that would work as well, the foregoing schedule being illustrative.

A further condition in which the adaptive gain may be altered from an otherwise usual gain value occurs when the adaptive filter \hat{h} is confronted with a far-end signal that is narrow band, i.e. comprised of a few sinusoids. In such a scenario, there are an infinite number of equally optimal solutions that the LMS adaptation scheme can find. Thus it is quite unlikely that the resulting cancellation solution \hat{h} will properly identify (i.e. mirror) the channel echo response h . Such a situation is referred to as under-exciting the channel, in that the signal only provides information about the channel response at a few frequencies. The echo canceller 25 attempts to determine the existence of this condition at step 220.

Consider a situation where the far-end signal varies between periods in which a narrow band signal is transmitted and wide band signal is transmitted. During the wide band signal periods, the \hat{h} filter should adapt to reflect the impulse response of the channel. However, when the narrow band signal transmission period begins, the \hat{h} filter may readapt to focus on canceling the echo path distortion only at the frequencies present in the narrow band signal. Optimizing a solution at just a few frequencies is likely to give a different solution than was found during transmission of the wide band signal. As a result, any worthwhile adaptation channel information gained during wide band transmission periods is lost and the \hat{h} filter requires another period of adaptation once the wide band signal returns.

When the far-end signal is narrow band, the adaptation can and should be slowed considerably, which should discourage the tendency of the coefficients to diverge. Specifically, when a narrow band signal is detected, a may be upper-bounded by either 0.25 or .125. This operation is illustrated at step 225.

Narrow band signal detection may be implemented using a low order predictive filter (e.g., a fourth order predictive filter). Preferably, this filter is implemented in software executed by one or more digital signal processors used in the echo canceller system 25. If it is able to achieve a prediction gain of at least 3 to 6 dB (user defined), then it is assumed that the received signal is a narrow band signal.

An amplitude threshold for the far-end signal is also preferably employed in determining the existence of a narrow band signal. If the far-end power is greater than -40 dBm, the current far-end sample is sent to the low order predictive filter, which determines whether or not the far-end signal is narrow band. If the far-end power is less than -40 dBm, the predictive filter is re-initialized to zero.

A further scenario in which it is desirable to alter the gain of the adaptive filter \hat{h} is when the echo path response is non-linear. The presence of nonlinearities in the echo path encourages constant minor changes in the coefficients \hat{h} in order to find short-term optimal cancellation solutions. The detection of non-linearity of the echo path response preferably proceeds in the manner set forth in U.S.S.N. ___, titled ___, filed on even date herewith. The presence of a non-linear echo path is determined that step 230.

In a non-linear echo path scenario, it is desirable to choose the adaptive gain constant a large enough that \hat{h} can track these short-term best solutions. However, choosing $a = 1$ may be suboptimal in most non-linear scenarios. This is due to the fact that the gain is too large and, thus, short-term solutions are "overshot" by the aggressive adaptation effort. Accordingly, as shown at step 235, choosing a gain lower than 1 is preferable. Choosing $a = 0.25$ was found to be the best trade off between tracking and overshooting short term optimal solutions. The gain constant a may be further reduced if large background noise is measured, as

discussed above.

A still further scenario in which the adaptive gain may be varied relates to the convergence period of the adaptive filter \hat{h} . As noted above, a large gain constant α is desired during convergence periods while a smaller α is desired in steady state conditions after the filter has converged. In other words, there seems little lost and perhaps some potential gain to reduce α after an initial period of convergence is completed. This appears to be especially valuable if the long-term performance is found to be substandard.

In view of the foregoing, the echo canceller 25 may implement a reduced gain mode in which an upper bound for the gain constant α is set at a lower value than 1 (e.g., at either .25 or .125). This mode is detected at step 240 and is entered at step 245 if the ERLE remains below a predetermined threshold value (e.g., either 6dB or 3dB) after a predetermined period of adaptation. The adaptation time is preferably selected as a value between 100 to 300 msec. This amount of time will generally prevent the echo canceller 25 from entering the reduced gain mode during convergence periods. The reduced gain mode may optionally be exited if the steady state ERLE increases above a certain threshold.

If the echo canceller does not enter the reduced gain mode at step 240, the gain constant α is preferably set or reset to a predetermined value. This operation is illustrated at step 250, where the gain constant $\alpha = 1$.

As discussed above, certain conditions can result in a lowering of the value of E_{max} . This is described in connection with Figure 4 above. Such an operation effectively results in a "soft reset" of the transfer operations. When this occurs, it may be desirable to clear all or some of the modes defined above and set the gain a to an initial value, for example, 1. These operations are illustrated at steps 300 and 305.

Separate and apart from the foregoing adjustments of the gain constant a , the present inventors have recognized that it may be advantageous to adapt a subset of the coefficients of the \hat{h} filter with a higher gain and the remaining coefficients with a smaller gain. To understand the motivations for doing this, consider a scenario in which the echo canceller 25 must converge to a linear echo-path. Since some flat-delay is to be expected, the span of time covered by the coefficients of the \hat{h} filter should be larger than the expected duration of the echo-path response. As a result, several of the taps (and in many cases, the majority of the taps) of the \hat{h} filter will have an expected value of zero to model the flat-delay while a small subset of "significant" taps will need to adjust very quickly in order to model the linear echo-path response.

In such a case, the convergence time is reduced when the "significant" taps are adapted separately from the smaller "flat-delay" taps. To see this, note from the Normalized LMS set forth above, that the adaptation gain increases as the

number of coefficients increase. If the significant taps are adapted separately, they will converge more quickly due to the fact that the adaptation process is directed to a fewer number of taps. Further, adaptation noise from all the flat-delay taps is minimized when they are adapted separately from the significant taps using a smaller gain $a' < a$, where a' is the gain for the flat delay taps and a is the gain for the significant taps. Thus, splitting may be helpful in steady-state if there is significant background noise.

Figure 7 illustrates an \hat{h} filter tap energy distribution for a typical linear echo path. The echo canceller 25 divides the taps into window sections, each window section preferably having the same number of taps. The echo canceller 25 then proceeds to determine which of the windows has the largest amount of energy disposed therein. The windows having the largest amount of energy are tagged as being more significant than other windows. The adaptive coefficients of the relatively few tagged windows are adapted separate from adaptive coefficients of the larger number of low energy, non-tagged windows. This naturally results in faster convergence of the coefficients of the tagged windows compared to convergence of the coefficients of the tagged windows in a non-split scenario. Still further, the adaptive coefficients of the tagged windows may be adapted using a higher gain constant a than the gain constant a' used to adapt the lower energy windows.

In accordance with a more specific embodiment of the split adaptation process, the echo processor 25 tags W sections of 8 contiguous taps each, where W is the number in msec of the echo response, not including flat-delay. Once chosen, these $W*8$ taps, or alternately W sections, are collectively called the in-window taps. Taps are considered for tagging as windowed-taps in contiguous blocks of 8 taps, representing 1 msec.

The $W*8$ taps are approximately the W sections which have the greatest energy. The tags are placed on the windows iteratively, that is, once one section is tagged, a new search is conducted to find the next tagged section. As an extension to the above, in order to encourage that these W tags will lump together into not more than a few larger sections, which is in often desired, two steps are taken. First, a section which immediately follows a tagged section is biased in the large-energy search by an additive or multiplicative constant, thus making it more likely to be chosen. Second, when a section is tagged due to the above search, one, two or more adjacent untagged sections are also tagged.

There are some scenarios in which it is undesirable to assume that some taps are more significant than others. Two examples are non-linear echo path scenarios, and narrow bandwidth scenarios (e.g., data calls having narrow bandwidths of the signaling frequencies).

Echo cancellation on non-linear paths is accomplished by finding short-term

minimizations of the time-varying performance surface. Echo cancellers for narrow bandwidth data calls need not properly identify the echo impulse response in order to be effective. In these two cases, no subset of the taps should be assumed more or less significant, and thus splitting gives non-optimal results. Thus splitting should be suppressed for non-linear calls and data calls.

As will be readily recognized, the echo canceller of the present invention may be implemented in a wide range of manners. Preferably, the echo canceller system is implemented using one or more digital signal processors to carry out the filter and transfer operations. Digital-to-analog conversions of various signals are carried out in accordance with known techniques for use by the digital signal processors.

Figure 8 illustrates one embodiment of an echo canceller system, shown generally at 700, that maybe used to cancel echoes in multi-channel communication transmissions. As illustrated, the system 700 includes an input 705 that is connected to receive a multi-channel communications data, such as a T1 transmission. A central controller 710 deinterleaves the various channels of the transmission and provides them to respective convolution processors 715 over a data bus 720. It is within the convolution processors 715 that a majority of the foregoing operations take place. Each convolution processor 715 is designed to process at least one channel of the transmission at line 730. After each convolution processor 715 has

processed its respective channel(s), the resulting data is placed on the data bus 720. The central controller 710 multiplexes the data into the proper multichannel format (e.g., T1) for retransmission at line 735. User interface 740 is provided to set various user programmable parameters of the system

Numerous modifications may be made to the foregoing system without departing from the basic teachings thereof. Although the present invention has been described in substantial detail with reference to one or more specific embodiments, those of skill in the art will recognize that changes may be made thereto without departing from the scope and spirit of the invention as set forth in the appended claims.